

The manipulation of chemical reactions: probing the limits of interventionism

Georgie Statham¹

Received: 5 November 2015 / Accepted: 9 July 2016

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract I apply James Woodward's interventionist theory of causation to organic chemistry, modelling three different ways that chemists are able to manipulate the reaction conditions in order to control the outcome of a reaction. These consist in manipulations to the reaction kinetics, thermodynamics, and whether the kinetics or thermodynamics predominates. It is possible to construct interventionist causal models of all of these kinds of manipulation, and therefore to account for them using Woodward's theory. However, I show that there is an alternate, more illuminating way of thinking about the third kind of reaction control, according to which chemists are thought of as manipulating which causal system is instantiated. I show that our ability to manipulate which system is instantiated is an important part of our ability to control the world, as is therefore especially relevant to interventionism. Thus, considering examples from organic chemistry leads to the identification of an important extension to Woodward's theory. Finally, this investigation into reaction control in organic chemistry also has a more general implication: it suggests that interventionism results in a version of pragmatism about causation.

Keywords Causation · Interventionism · Woodward · Chemistry

1 Introduction

To set up a reaction, chemists typically take a flask, add solvent and starting materials, then leave the resulting mixture to stir at a given temperature for a certain amount of time. The overall aim is to produce a particular chemical species in as high a yield as

✉ Georgie Statham
gms47@cam.ac.uk

¹ Trinity College, The University of Cambridge, Cambridge CB2 1TQ, UK

possible; it is by varying the starting materials, solvent, temperature, etc., that chemists are able to synthesise different chemical species.

This procedure looks very much like interventionist causation: a chemical reaction is a causal process in which one set of chemical species is converted into another set of chemical species, and chemists are able to intervene on this process, influencing the product of the reaction by manipulating a range of variables. At a cursory glance, chemical reactions therefore seem well suited for interventionist causal modelling, and several authors have proposed that the causal claims made (or explanations given) in organic chemistry should be understood in terms of this theory (Goodwin 2012, 2013; Ramsey 2008; Statham 2016).

Chemists are able to manipulate the outcome of a reaction by either changing the starting materials, or manipulating the reaction conditions. After showing that interventionism can account for the former, I focus on the latter, a form of reaction control that has not previously been discussed in the philosophical literature.¹ Chemists use the reaction conditions to manipulate both the kinetics and thermodynamics of chemical reactions, as well as whether the kinetic or thermodynamic product is favoured. Modelling these forms of manipulation provides a rigorous test of the applicability of Woodward's interventionism to organic chemistry.

I show that all three forms of manipulation can be represented within Woodward's framework. However, there is also an alternative way of representing manipulations that control whether the kinetic or thermodynamic product is formed, namely as interventions on the causal system itself. This kind of intervention can't be captured by interventionist causal models, and therefore can't be accounted for by the interventionist theory. I show that interventionism needs to be able to account for interventions on the causal system, because this kind of manipulation is central to our ability to control the world. This category therefore represents a form of causal reasoning that has previously been overlooked by Woodward's theory, but that needs to be incorporated—that is, it represents a limitation of interventionism.

Finally, I use this investigation into the control of organic reactions to address another question that is central to interventionism: what is an appropriate causal model? The theory of organic chemistry is shaped by its main goal, namely synthesis. The causal models discussed here are therefore also shaped by this goal. This suggests that what counts as an appropriate causal model is partly determined by the goal of a particular causal enquiry. Thus, interventionism leads to a form of pragmatism about causation.

2 Woodward's interventionism

According to Woodward, variable X is a cause of variable Y if and only if an intervention that changed the value of X would result in a change to the value of Y . The term 'intervention' is a technical notion: roughly, an intervention on X with respect to Y is a change in the value of X that affects Y (if at all) only via a pathway that

¹ For simplicity and clarity, I limit the discussion to organic chemistry. However, much of the following also applies to inorganic chemistry.

goes through X .² So, for example, smoking is a cause of cancer because if we were to (unethically) use a random controlled experiment to control the amount that a group of subjects smoked, we would discover that the amount smoked makes a difference to an individual's risk of cancer.

Interventionists represent causal systems using causal models. These consist of a set of exogenous and endogenous variables, structural equations, and background variables. The structural equation for each endogenous variable (i.e., each variable that is an effect of other variables in the graph) expresses the value of that variable as a function of its direct causes. Diagrammatically, the variables are connected by arrows (or directed edges), which indicate that a direct causal relationship obtains between two variables. The background variables are held fixed, and can be thought of as representing the context in which the system of structural equations applies.³

Woodward's account entails that it is possible to represent every (actual) causal relation as part of a causal model with the above characteristics. In order to use organic chemistry to test his theory, it is therefore necessary to construct models of causal systems in this field. As noted above, I have chosen to model three related systems representing the various ways that chemists are able to use the reaction conditions to influence the outcome of a reaction. This approach is unusual: philosophers typically use thought experiments and/or isolated examples to test theories. By choosing related examples, my strategy focuses attention on the connections between models, as well as the individual models themselves. It is this unusual focus that brings into view the possibility of intervening on the causal system itself.

3 Controlling organic reactions

The most common—and most versatile—way of creating different chemical species is to use different starting materials. However, because any set of starting materials will interact differently in different conditions, chemists also need to carefully control the reaction conditions. In fact, changing the reaction conditions will often result in the formation of different products. This therefore represents a second way of controlling the outcome of a reaction. I briefly discuss the former kind of control, before focusing on the latter.

3.1 Manipulating the starting materials

We know that (in most cases) changing the identity of the starting materials makes a difference to the identity of the product. According to interventionism, this means that the variable {identity of the starting materials} (SM) is causally relevant to the variable {identity of the product} (P). For example, using chloride rather than bromide in a nucleophilic substitution reaction will cause the product to be chlorinated rather than brominated.

² For a more precise definition of the notion of an intervention, see Woodward (2003, pp. 98–99).

³ Woodward says that X is a cause of Y if and only if an intervention that changed the value of X would result in a change to the value of Y , in some background circumstances.

It is not much use knowing that changing the starting materials will make a difference to the product of a reaction unless you have some idea of what effect any given change will have. Unsurprisingly, much of the theory of organic chemistry is geared towards making such predictions. In earlier work in the philosophy of organic chemistry, William Goodwin shows that organic chemists have developed a number of structural and electronic concepts that allow chemists to make qualitative generalisations that hold between kinds of chemical species (2003, 2012). For example, the concept of a *nucleophile* allows chemists to predict that chloride, bromide, and hydroxide can all be added to a carbon centre using a nucleophilic substitution reaction, whereas the nitronium ion (an electrophile) cannot. With this and other structural and electronic concepts in hand, organic chemists are able to categorise chemical species and make predictions about the reactions that whole classes of species will undergo (Goodwin 2012, pp. 437–438). In terms of interventionism, these concepts allow chemists to determine which sets of starting materials can be used to make which products—that is, which values of *SM* correspond to which values of *P*.

Manipulations involving changes to the starting materials can therefore be well accounted for by interventionism. My focus, however, is the second form of reaction control introduced above, in which chemists hold the starting materials fixed and manipulate the reaction conditions. The reaction conditions can influence the outcome of a reaction in two different ways: by affecting the reaction *kinetics* or *thermodynamics*. I now introduce these two concepts.

3.2 Kinetics and thermodynamics

The term ‘kinetics’ relates to the rate at which a reaction occurs—that is, the rate at which the starting materials are used up and the products form. ‘Thermodynamics’, on the other hand, involves the relative stabilities of the reactants and products, and therefore which product is energetically favoured.⁴ Corresponding to these two terms, chemists define two kinds of product. The *kinetic product* of a reaction is the species that forms fastest (i.e., the product with the highest rate of reaction). The *thermodynamic product* is the species that is lowest in energy, and therefore most stable.⁵

If a reaction has two (or more) possible products that are *not* interconvertible under the reaction conditions, the major product will always be the kinetic product. This species forms first, and is stable under the reaction conditions. Thus, once it has formed, the reaction is over. Thermodynamics becomes relevant when two possible products (or the products and starting materials) *are* interconvertible. In this case, the two products eventually reach an equilibrium that favours the thermodynamic product.

⁴ For an introduction to kinetics and thermodynamics in organic chemistry, see Clayden et al. (2012).

⁵ It is worth noting that reactions do not always have both a kinetic and a thermodynamic product. In order to count as a possible product at all, a chemical species must be stable enough to isolate, and must actually be able to form under some realistic set of reaction conditions. Thus, if a reaction is such that it never reaches equilibrium, it won’t have a thermodynamic product, and if it reaches equilibrium extremely quickly (even at low temperatures), it won’t have a kinetic product.

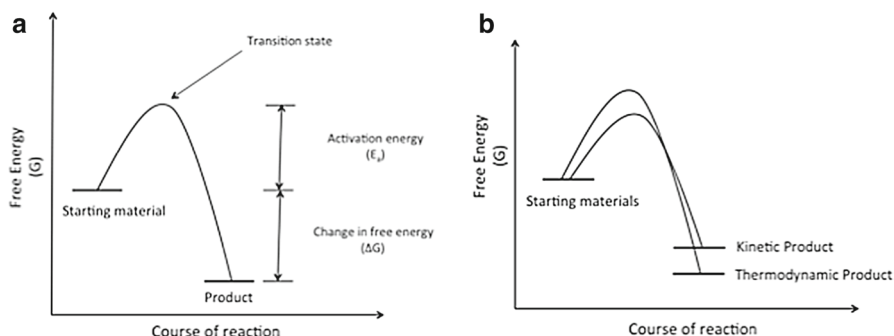


Fig. 1 **a** Example energy profile diagram. **b** Kinetic versus thermodynamic products

Kinetics and thermodynamics can be most clearly illustrated using a reaction profile diagram. These plot the energy of the reacting species against the geometrical path of the reaction, as shown in Fig. 1a.

The difference in energy between the starting materials and the products is the *free energy of reaction* (ΔG). The lower the energy of the products relative to the starting materials (i.e., the more negative ΔG), the more thermodynamically favoured the reaction. The thermodynamic product is therefore the product with the most negative value of ΔG , as shown in Fig. 1b.

The highest point on the reaction profile diagram represents the *transition state*—a (relatively) high energy intermediate that the reaction must pass through in order to form the products. The difference in energy between the starting materials and the transition state thus represents an energy barrier, called the *activation energy* (E_a), that must be overcome if the reaction is to proceed.⁶ The lower the activation energy, the faster the reaction. The kinetic product is therefore the product with the lowest activation energy. Again, this is shown in Fig. 1b.

Corresponding to the terms ‘kinetic product’ and ‘thermodynamic product’, respectively, chemists divide reactions into two classes. Reactions under *kinetic control* are those in which the major product is the kinetic product; in reactions under *thermodynamic control* the major product is the thermodynamic product. This distinction is important, because reactions that are under kinetic control are governed by (and therefore can be described by) kinetic laws, whereas reactions under thermodynamic control are governed by (and described by) quite different thermodynamic laws.

3.3 Manipulating the reaction conditions

Because the reaction kinetics and thermodynamics are governed by different physical laws, they can be manipulated by changing the reaction conditions in different ways. By manipulating (i) the kinetics or (ii) the thermodynamics, respectively, chemists can often influence which species is the kinetic product and which is the thermodynamic

⁶ There may be more than one transition state, in which case the activation energy is the difference between the starting materials and the highest energy transition state.

product. In some cases, it is also possible to control (iii) whether the kinetic or thermodynamic product is favoured. Each of these three scenarios requires intervening on different variables, according to different (usually only qualitatively specified) rules.

For example, consider a set of starting materials (A and B) that undergo two competing reactions, the first of which yields products C and D, and the second of which yields E and F. Assume that C is the desired product. D is therefore a byproduct; E and F are the species we want to prevent from forming.

Desired reaction: $A + B \rightarrow C + D$

Undesired alternative: $A + B \rightarrow E + F$

C and E could be (i) two possible kinetic products; (ii) two possible thermodynamic products; or (iii) C could be the kinetic product and E the thermodynamic product, or vice versa. These three possibilities correspond to three scenarios discussed above. In (i) chemists are able to manipulate which product is the kinetic product; in (ii) it is possible to manipulate which product is the thermodynamic product (or in which direction an equilibrium is pushed); and finally, in (iii) chemists can influence whether the kinetic or thermodynamic product is favoured. I refer to these scenarios as ‘kinetic manipulation’, ‘thermodynamic manipulation’, and ‘kinetic versus thermodynamic manipulation’, respectively.⁷

The above three-way classification is an alternative to the two-way classification used by chemists (in which reactions are classified as being under kinetic or thermodynamic control). This alternative classification makes causal modelling possible, because within each category, the same set of variables are causally relevant to the outcome of the reaction.

Organic chemistry textbooks don’t attempt to model the various ways that chemists manipulate the outcome of reactions. Instead, these textbooks tend to include a brief introduction to the relevant physical chemistry (i.e., kinetics and thermodynamics), then a series of examples that illustrate the various ways of controlling reaction outcomes.⁸ Each example can be categorised as involving kinetic, thermodynamic, or kinetic versus thermodynamic manipulation.⁹ These three categories are separate, in the sense that for a given reaction, chemists only carry out one of the three kinds of manipulation.

The point of modelling these various forms of reaction control is not to provide a tool for organic chemists. Rather, the models included in this paper are of value for two other reasons. First, they make explicit some of the theory of organic chemistry, and could thus be used as a case study against which to assess theories in the philosophy

⁷ The taxonomy of reactions described above is simplified in at least two ways. First, there will generally be more than two possible products of a reaction, and thus more than one undesired alternative. Second, chemical reactions do not always have the form $A + B \rightarrow C + D$. The models described below can be adapted to account for reactions with both more possible products and with formulae of different forms (e.g., $A \rightarrow C + D$, $A + B \rightarrow C$, etc.). However, trying to include all these forms into a single model would result in a graph that was too complicated to realistically illustrate.

⁸ For example, see Clayden et al. (2012).

⁹ Each example can also be categorised as under kinetic or thermodynamic control, as follows: Reactions that are under kinetic or thermodynamic manipulation are also under kinetic or thermodynamic control, respectively. In the case of kinetic versus thermodynamic manipulation, precisely what is being manipulated is whether the reaction is under kinetic or thermodynamic control.

of science and the philosophy of causation. Second, and more importantly for current purposes, it is necessary to model the causal reasoning used by organic chemists in order to use this discipline as a test case for Woodward's interventionism.

4 Kinetic manipulation

Organic chemists use kinetic manipulation to control which product is formed in situations in which there are two (or more) possible kinetic products. In this case, the overall product is the species that forms fastest—that is, the kinetic product. The aim is therefore to maximise the rate of the desired reaction (relative to the undesired reaction), in order to ensure that the yield is as high as possible. Using a selective catalyst is one way of achieving this, as illustrated in the following example.

Figure 2 shows the key step in the synthesis of a conduritol derivative—a chemical species that is of synthetic interest because some derivatives have antifeedant, antibiotic, antileukemic, and growth-regulating properties. In the step illustrated, the desired reaction is a *ring closing metathesis*, in which an eight-membered carbon chain is converted into a six-membered ring: a cyclohexene. This reaction only works with specific catalysts, so when the synthesis was developed, finding an appropriate catalyst was crucial (Ackermann et al. 2000).

The starting material for the metathesis reaction is the diene, **1**. A diene is a species containing two carbon-carbon double bonds—it is these double bonds that bind to the catalyst, reducing the activation energy of the metathesis reaction. Ackermann et al. report that when **1** is treated with a specific ruthenium catalyst, it undergoes the ring closing metathesis, as hoped, to give the cyclohexene **2**. However, when treated with a different, molybdenum-based catalyst, **1** undergoes quite a different reaction. In this case the starting material isomerises, yielding the acyclic species **3**.¹⁰

In both cases, the catalyst works by making a lower energy pathway available. This results in a lower energy transition state, thus reducing the activation energy and increasing the rate of the relevant reaction. However, the two catalysts make different lower energy pathways available, resulting in different products. This example therefore illustrates one way in which changing the reaction conditions can alter which of two possible products is the kinetic product.

Any intervention that changes the relative rates of two competing reactions that are under kinetic control will result in a change in the product distribution. We can systematise this by considering some simple kinetic laws. These laws apply to reactions that are either irreversible, or in which the rate of the forward reaction is much greater than the rate of the reverse reaction. In these situations, the rate of reaction is given by Eq. 1:

$$rate = k[A][B] \quad (1)$$

¹⁰ A note about the notation used in chemical reaction diagrams: carbon atoms are represented by line vertices or ends. Hydrogen atoms bonded to carbon are not directly represented, but are assumed to occupy the remaining bonding sites (each carbon forms four bonds). Atoms of other elements are represented by their chemical symbols. The wedge shaped bonds represent relevant three dimensional information: a solid wedge represents a bond coming out of the page, and a dashed wedge represents a bond going into the page.




Fig. 2 The effect of two different catalysts on the reaction of a substituted diene

Here, k is the rate constant for the reaction, and $[A]$ and $[B]$ are the concentrations of the species involved in the rate-determining step of the reaction.¹¹

The rate constant for any particular reaction is given by the Arrhenius equation:

$$k = Ae^{-Ea/RT} \quad (2)$$

Here, A is a quantity known as the pre-exponential factor (this is constant for a given reaction), Ea is the activation energy, R is the gas constant, and T is the temperature.

A pair of molecules will react when they collide with enough energy to overcome the activation energy. So, more descriptively, k is the number of collisions that result in reaction per second, A is the number of collisions per second (at the right orientation and at a given concentration), and $e^{-Ea/RT}$ is the probability that any given collision will result in reaction.

Combining these equations, we get Eq. 3:

$$rate = Ae^{-Ea/RT}[A][B] \quad (3)$$

¹¹ Equation 1 describes the rate of a second-order reaction, in which the rate-determining step involves the collision of a molecule of A with a molecule of B. First- and third-order rate equations are also common, but will not be discussed here.

This equation shows that the activation energy, the temperature, and the concentrations of the species involved in the rate-determining step all influence the rate of a reaction, and can therefore be used in kinetic manipulation.

$[A]$ and $[B]$ represent the instantaneous concentrations of species A and B. When setting up a reaction, chemists are generally only able to manipulate the initial concentration of the reactants, $[A_o]$ and $[B_o]$. However, if a reaction is under kinetic control, the major product will be the species with the highest initial rate of reaction. This means that we can replace the instantaneous concentrations in Eq. 3 with initial concentrations, to give:

$$rate = Ae^{-Ea/RT} [A_o][B_o] \quad (4)$$

Although Eq. 4 only applies quantitatively to the initial stages of a reaction, it can be used to qualitatively predict the relative rates of two reactions throughout the course of the reaction.

The temperature of the reaction and the initial concentrations of A and B can all be manipulated directly. However, the activation energy can only be manipulated indirectly, most typically by using a catalyst or changing the solvent.

This completes the information required to construct a causal model of the variables used in kinetic manipulation, given in Fig. 3. This model is intended to represent the scenario in which a chemist wants to maximise the yield of the kinetic product, while preventing the formation of an alternative kinetic product. Reaction 1 is thus $A + B \rightarrow C + D$; reaction 2 is $A + B \rightarrow E + F$. The concentrations are initial concentrations, and the reaction rates are those that obtain during the first 5–10 % of the reaction.

Variables

Ca: the identity of the catalyst (or the absence of a catalyst)

So: the identity of the solvent

SM: the identities of A and B

*Ea*₁: the activation energy of reaction 1

*Ea*₂: the activation energy of reaction 2

$[A_o]$: the initial concentration of A

$[B_o]$: the initial concentration of B

T: the temperature

*R*₁: the initial rate of reaction 1

*R*₂: the initial rate of reaction 2

KP: the identity of the kinetic product

Structural equations


$Ea_1 = f(Ca, So)$, given *SM*

$Ea_2 = f(Ca, So)$, given *SM*

$R_1 = A_1 \cdot e^{-Ea_1/RT} \cdot [A_o] \cdot [B_o]$ (where A_1 is the pre-exponential factor for reaction 1, a constant)¹²

¹² The structural equations for R_1 and R_2 assume that both these reactions have second-order kinetics. Different equations will apply if the reactions are first- or third-order.

Fig. 3 Variables controlling the product distribution of reactions under kinetic manipulation



$$R_2 = A_2 \cdot e^{-Ea_2/RT} \cdot [A_o] \cdot [B_o] \text{ (where } A_2 \text{ is the pre-exponential factor for reaction 2)}$$

$$KP = C \text{ if } R_1 > R_2; E \text{ if } R_2 > R_1$$

The causal model in Fig. 3 (as well as the other causal models that will be discussed) is intended to be understood in the causal modelling tradition, as discussed in Sect. 2.

Figure 3 shows that the catalyst and the solvent both influence the activation energy of a reaction. The activation energy, the temperature, and the concentration of the starting materials then affect the rate of each reaction, and the relative rates of the two reactions determine which is the kinetic product.¹³

The arrows between *SM* and both *Ea*₁ and *Ea*₂ are dashed to represent the fact that, although the starting materials are causally relevant to the activation energies, these are held fixed in the causal system being depicted, and are thus not possible sites of intervention.¹⁴

Organic chemists must have implicit knowledge of the causal structure represented by the model in Fig. 3 (as well as knowledge of the mechanisms of the relevant reaction kinds) in order to fully understand (and make) the causal explanations and predictions used in this field. For example, consider the reactions illustrated in Fig. 2. With respect to these reactions, it is true to claim that:

Using the ruthenium catalyst (rather than the molybdenum catalyst) on the diene will promote the ring closing reaction (rather than the isomerisation).¹⁵

¹³ Although there is not space for a detailed discussion here, it can be shown that the variables included in the above model are all distinct, capable of being intervened on in the right way, and that the system of equations is modular.

¹⁴ Of course, *SM* is not the only variable that is causally relevant to the outcome of reactions under kinetic control, but that is held fixed in the above model. There are numerous other variables that comprise the background and are therefore held fixed: for example, the fact that the reaction is carried out in an inert flask. It is a feature of all interventionist models that the variable choice is partly pragmatic. See Woodward (2003, pp. 86–91); Hitchcock (2001).

¹⁵ The above sentence is an example of a contrastive causal judgement—that is, a judgement of the form: *C* rather than *C** is a cause of *E* rather than *E**. It is an implication of Woodward's interventionism that causal judgements are contrastive in form (2003, pp. 145–146). For other contrastive accounts of causation, see Schaffer (2005, 2013), Maslen (2004) and Hitchcock (1996).

In order to give a full explanation of this claim, you need to know that catalysis decreases the activation energy of a reaction, which increases the rate of the catalysed reaction, and thus favours the reaction in question. That is, you need to understand part of the causal structure represented by the model in Fig. 3.¹⁶ This model can therefore be considered as a representation of part of the theory of organic chemistry.

The fact that it is possible to construct a causal model of kinetic manipulation shows that this kind of reaction control can be nicely accounted for by Woodward's interventionism. In the next section, I show that thermodynamic manipulation can also be accounted for by this theory.

5 Thermodynamic manipulation

Chemists use thermodynamic manipulation to control which of two possible thermodynamic products is favoured, or which of two species at equilibrium is the major product. Recall that the thermodynamic product is the product that is lowest in energy, and that is therefore favoured when the reaction reaches thermodynamic equilibrium. Once a reaction has reached equilibrium, it is possible to drive the reaction in either direction by adding or removing components of the equilibrium. This can be illustrated using Fisher esterification, shown in Fig. 4.

Fisher esterification is a common reaction, used to make an ester from a carboxylic acid and an alcohol, using an acid catalyst. The R and R' groups in the reaction diagram represent carbon chains of any length: the reaction can therefore be used to convert a wide range of carboxylic acids into an even wider range of esters.

The Fisher esterification reaction is under thermodynamic control, with an equilibrium position that favours the starting materials—that is, the carboxylic acid is also the thermodynamic product. However, the reaction can be driven to completion by using the alcohol as the solvent (i.e., ensuring that the concentration of the alcohol is very high), and/or removing water as it forms. The reason for this will become clearer if we look at some of the theory of reaction thermodynamics.

The equilibrium position of a reaction is described by the equilibrium constant (K). K is constant for a given reaction at a given temperature; its value is a function of the free energy change for a reaction (ΔG) and the temperature (T), as shown in Eqs. 5 and 6. (R is the gas constant.)

$$\Delta G = -RT \ln K \quad (5)$$

$$\therefore \ln K = -\Delta G / RT \quad (6)$$

K is a ratio of the concentrations of the products and the starting materials at equilibrium. For example, for the reaction: $A + B \rightleftharpoons C + D$, the equilibrium constant is given by Eq. 7:

$$K = \frac{[C][D]}{[A][B]} \quad (7)$$

¹⁶ I take it that what it takes to give a full explanation is context sensitive. The 'full explanation' referred to above is therefore a full explanation within the context of organic chemistry.




Fig. 4 Fisher esterification

If K is greater than one, the thermodynamic product is C , and if K is less than one, the thermodynamic product is A . The identity of the thermodynamic product can therefore be determined from the value of K .

However, the thermodynamic product is not always the major product of a reaction under thermodynamic control. Since the equilibrium constant is constant for a given reaction at a given temperature, the concentration of any species involved in the equilibrium state can be altered by changing the concentrations of the other species. It is therefore possible to push the reaction in either direction (i.e., towards or away from the thermodynamic product) by adding or removing product or starting material. It is this process that is exploited in Fisher esterification: by either using an excess of the alcohol (one of the starting materials) or removing water (the byproduct), the reaction is driven towards formation of the ester.

Given the above information, Eqs. 6 and 7 can be used to construct a causal model of the factors influencing the product of a reaction under thermodynamic manipulation. This model is given in Fig. 5.¹⁷ Note that this model applies only to reactions that reach an equilibrium (either between products and starting material, or between two possible products). Again, reaction 1 is $A + B \rightarrow C + D$, and reaction 2 is $A + B \rightarrow E + F$.¹⁸

Variables

SM : the identities of A and B

ΔG_1 : the change in free energy for reaction 1

ΔG_2 : the change in free energy for reaction 2

T : the temperature

K : the equilibrium constant between the products of reaction 1 and the products of reaction 2 = $\frac{[C][D]}{[E][F]}$ (where these are the concentrations at equilibrium)

TP : the identity of the thermodynamic product

C_r : the amount of C removed from the reaction

D_r : the amount of D removed from the reaction

E_r : the amount of E removed from the reaction

F_r : the amount of F removed from the reaction

P : the identity of the major product

¹⁷ The free energy of a reaction depends on the change in enthalpy (the energy produced by the formation of bonds and used in the breaking of bonds), the change in entropy (the relative disorder of a reaction solution before and after the reaction), and the temperature, according to the equation: $\Delta G = \Delta H - T\Delta S$. In the interests of simplicity, I do not discuss this detail in the main text.

¹⁸ Notice that in Fisher esterification, the starting materials are also the alternative products—that is $A = E$ and $B = F$.




Fig. 5 Variables controlling the product distribution of reactions under thermodynamic manipulation

Structural equations

$$\Delta G_1 = f(T), \text{ given } SM$$

$$\Delta G_2 = f(T), \text{ given } SM$$

$$K = e^{-\frac{(\Delta G_1 - \Delta G_2)}{RT}} \text{ }^{19}$$

$$P = f(K, C_r, D_r, E_r, F_r)$$

This model shows that changing either the temperature or the free energy of a reaction will make a difference to the equilibrium constant, and therefore the identity of the thermodynamic product. Additionally, by adding or removing any of the equilibrium components, it is possible to push the reaction in the direction of one of the products.

The directed arrow between the variables K and TP is dashed, because these two variables fail to be *distinct*: the value of K entirely determines the identity of the thermodynamic product (which is C if $K > 1$, and E if $K < 1$).²⁰ For this reason, the relation that holds between K and TP should be thought of as determination, not causation.²¹

Like the model of kinetic manipulation discussed above, the causal structure illustrated in the model of thermodynamic manipulation must be understood in order to fully understand some of the predictions made and explanations given in organic chemistry. For example, consider the following statement, made with respect to the Fisher esterification reaction:

Removing water will favour the formation of the ester (rather than the carboxylic acid).

¹⁹ K is the equilibrium position between the products of reaction 1 and the products of reaction 2—that is, for the reaction $C + D \rightarrow E + F$. The value of ΔG for this reaction is therefore $(\Delta G_1 - \Delta G_2)$, and $K = \frac{[E][F]}{[C][D]}$.

²⁰ An important constraint on the variables included in an interventionist causal model is that these variables themselves must be *distinct*, in the sense articulated by David Lewis (1986)—that is, one variable taking a particular value should not logically imply that a second value takes a particular value. See Hitchcock (2001, 2007).

²¹ The other variables in the model are all distinct in the sense required for interventionism: there is no logical or conceptual connection between the values that they take. The structural equations are also modular.

In order to give a full explanation of this claim, you need to know that the starting materials and products are in equilibrium, that water is a byproduct of the reaction, and therefore that removing water will push the reaction in the direction of the products. Thus, again, the model in Fig. 5 represents part of the theory of organic chemistry. The possibility of constructing such a model shows that thermodynamic manipulation, like kinetic manipulation, can be accounted for by Woodward's interventionism.

In the next section, I argue that kinetic versus thermodynamic manipulation can also be modelled using an interventionist graph. However, a more theoretically illuminating way of thinking about these manipulations is as interventions on the causal system itself. This alternative can't be represented by a causal model and therefore can't be accounted for by the interventionist theory.

6 Kinetic versus thermodynamic manipulation

Kinetic versus thermodynamic manipulation occurs when chemists control whether the kinetic or the thermodynamic product of a reaction is favoured—that is, whether the reaction is under kinetic or thermodynamic control. An example is the reaction of 1-phenylpropyne with HCl, illustrated in Fig. 6.

Phenylpropyne reacts with HCl to give two isomers of the corresponding chloroalkene. The kinetic product of this reaction is the E-alkene, the formation of which has a lower activation energy than that of the Z-alkene. However, the thermodynamic product is the Z-alkene, presumably because there is less steric hindrance between the phenyl and methyl groups in this isomer.

If the reaction is carried out at low temperatures for a short period of time, it is possible to isolate the E-alkene (the kinetic product). By using higher temperatures and longer reaction times, however, the reaction can be brought to thermodynamic equilibrium, at which point more than 95 % of the reaction mixture has been converted into the Z-alkene (the thermodynamic product).




Fig. 6 Reaction of 1-phenylpropyne with HCl

Fig. 7 Variables controlling the product distribution of a reaction under kinetic versus thermodynamic manipulation




Figure 7 represents the causal structure that appears to underlie kinetic versus thermodynamic manipulation in the reaction of phenylpropylene with HCl, and in general.

Variables

- KP*: the identity of the kinetic product
- RT*: the reaction time
- T*: the temperature
- TP*: the identity of the thermodynamic product
- P*: the identity of the overall product

Structural equation

$$P = f(RT, T), \text{ given } KP \text{ and } TP$$

When controlling whether the kinetic or thermodynamic product is isolated, chemists manipulate the reaction time and the temperature. It is assumed that the identity of the kinetic and thermodynamic products are unchanged by these manipulations. The arrows linking the variables *KP* and *TP* to *P* in Fig. 7 are therefore dashed to represent the fact that, although these variables are causally relevant to the overall product, they are assumed to be held fixed for the purposes of kinetic versus thermodynamic manipulation.

The model in Fig. 7 is one way of representing the control of organic reactions, which is consistent with interventionism. However, I now show that there is an alternative way of thinking about kinetic versus thermodynamic manipulation.²²

²² It may appear that the models of kinetic, thermodynamic, and kinetic versus thermodynamic manipulation (Figs. 3, 5, 7) can be combined to give an overall model of the ways that manipulating the reaction conditions can affect the outcome of a reaction. This is not in fact possible, however, because the three models represent different scenarios. Recall that the model of kinetic manipulation is intended to represent reactions in which the two most likely products are both possible kinetic products, the model thermodynamic manipulation is intended to represent reactions that have two possible thermodynamic products, and the model of kinetic versus thermodynamic manipulation represents reactions that have both a possible kinetic and thermodynamic product. This means that what is represented by the variables *KP* and *TP* is in fact subtly different in the different models. In the model of kinetic manipulation, *KP* represents the faster of two competing kinetic reactions, whereas in the model of kinetic versus thermodynamic manipulation, *KP* represents the product with the fastest rate of reaction, as opposed to the the most thermodynamically stable product. Similarly for *TP*.




Fig. 8 Reaction time vs concentration of kinetic and thermodynamic products, showing the time slices at which the kinetic and thermodynamic models apply

6.1 Kinetic versus thermodynamic manipulation as consisting in interventions on the causal system

We can arrive at an alternative way of thinking about the relationship between the variables KP , TP , RT , T , and P by considering the following: When a chemist manipulates whether the kinetic or thermodynamic product is favoured, she is really manipulating how far the reaction is allowed to progress towards equilibrium. This is illustrated in Fig. 8, a graph of reaction time versus concentration of the kinetic and thermodynamic products.

Reactions that are stopped at different stages can be described by different sets of laws. These different stopping points can therefore be thought of as instantiating different kinds of causal systems, represented by different causal models. This suggests that, in the case of kinetic versus thermodynamic manipulation, changing the reaction time and/or temperature can be thought of as intervening to influence *which kind of causal system is instantiated*. Consider the situation in which a chemist reduces the temperature and the reaction time in the reaction of 1-phenylpropyne and HCl. We can think of the chemist as manipulating the reaction conditions in order to ensure that kinetic laws govern the outcome of the reaction—and therefore that the causal system represented by the model of kinetic manipulation is instantiated, and the kinetic product is favoured.

On this interpretation, the causal structure of kinetic versus thermodynamic manipulation is represented in two stages. Stage one represents the use of temperature and reaction time to control which system is instantiated—that is, which set of laws (kinetic or thermodynamic) governs the outcome of the reaction, and therefore which of the kinetic or thermodynamic models applies. The relevant kinetic or thermodynamic model then comprises stage two.²³ Thus, the variables represented in stage one

²³ In most cases, the kinetic and thermodynamic models in stage two will only include one possible product, rather than two competing reactions.

determine whether the reaction is under kinetic or thermodynamic control; stage two illustrates the variables that affect the kinetics or thermodynamics themselves.

A significant advantage of this alternative interpretation of kinetic versus thermodynamic manipulation is that it unifies all the reactions that are under *kinetic control*, as well as all the reactions under *thermodynamic control*. The concepts ‘kinetic control’ and ‘thermodynamic control’ are central to the way that chemists think about reactions. Recall that chemists use these concepts to categorise reactions: any reaction in which the major product is the product that forms fastest is said to be under kinetic control, whereas reactions that reach equilibrium are said to be under thermodynamic control.

Now, the set of reactions under kinetic control includes all instances of kinetic manipulation, *plus* the instances of kinetic versus thermodynamic manipulation in which the kinetic product is favoured. Correspondingly, the set of reactions under thermodynamic control includes all cases of thermodynamic manipulation, as well as reactions under kinetic versus thermodynamic manipulation in which the major product ends up being the thermodynamic product. The idea that kinetic versus thermodynamic manipulation involves intervening on the causal system itself can therefore relate this form of manipulation to the concepts of kinetic and thermodynamic control: on this interpretation, all reactions under kinetic control are represented by the model of kinetic manipulation in stage two, whereas all reactions under thermodynamic control are represented by the model of thermodynamic manipulation. Thus, the advantage of this way of thinking about kinetic versus thermodynamic manipulation is that it illuminates the links between different causal systems—that is, between kinetic, thermodynamic and kinetic versus thermodynamic manipulation.

Which of the two interpretations is preferable depends on whether a chemist is concerned only with being able to carry out kinetic versus thermodynamic manipulation, or whether she wants to understand how this form of manipulation is connected to kinetic and thermodynamic manipulation. Thus, which model is appropriate depends a chemist’s explanatory goals.

In the next section, I show that the idea of being able to intervene on the causal system is compatible with Woodward’s interventionism; indeed, this category of manipulation is an important extension to the existing theory.

7 Interventions on the causal system

Woodward claims that the purpose of the concept of causation is to identify correlations that are exploitable for the purposes of manipulation and control (2003, pp. 9–12). In this respect, he takes the conclusion of Cartwright’s highly influential 1979 paper seriously: we need the concept of causation to distinguish between effective and ineffective strategies.

In order for it to be possible to exploit causal relationships to bring about desired ends, we need to see ourselves as able to freely intervene on the world—that is, we need to see (at least some) human actions as independent of their past causal histories. Woodward’s notion of an intervention has its origins in these ideas about manipulation and control. On his theory, a variable that can (in principle) be used

to intervene on the value of X is a variable that can (in principle) be used to set the value of X , independently of X 's causal history. Thus, interventions don't have to be human actions, but they retain the feature of actions that makes actions useful for manipulation and control. Free actions are therefore paradigmatic examples of interventions.

One consequence of centering a theory of causation around the above notion of an intervention is that the theory applies most naturally to 'small worlds'—that is, to systems that comprise only part of the world. Let's think about human actions again: these can only be external manipulations relative to a system that doesn't include the action itself, and its causal history. They must therefore be interventions on systems that comprise part of the universe, rather than the whole universe. Any model of such a system will have a restricted context of application—in interventionist terms, it will have a limited invariance range.²⁴

Consider the following example of a small world system: The effects of displacing a spring are usually modelled using Hooke's law, $F = -kX$. We can construct an interventionist model of a spring-mass system, showing that the displacement of a spring is proportional to the force applied. However, Hooke's law doesn't have universal application: for example, it is possible to apply so much force that the spring breaks. In this situation, neither Hooke's law itself, nor the corresponding interventionist model, applies. An intervention that sets the force higher than the breaking force of the spring therefore lies outside the invariance range of the original causal model.²⁵

The examples that are usually used to illustrate this phenomenon tend to focus on the fact that a causal model only represents a causal system within a certain context, or invariance range. However, if the focus is directed towards the *change* in causal model instead of the *limits* of a causal model, it becomes evident that part of causal reasoning must involve being aware of the invariance range of the causal model you are (either explicitly or implicitly) working with, and adopting a new model if the situation evolves outside this invariance range. When this happens, an intervention has changed the causal system—the small world you were using to reason with—itself.

When considering the forms of reaction control discussed in this paper, it becomes evident that it is commonplace for chemists to manipulate the causal system that is instantiated. Chemists ensure that reactions are set up such that they are either under kinetic or thermodynamic control—that is, they can be stopped within one of the two time slices highlighted in the graph in Fig. 8. Kinetic versus thermodynamic manipulation takes this a step further. Here, chemists are explicitly aware of being able to manipulate the reaction conditions such that the reaction is under either kinetic

²⁴ See Woodward (2007, pp. 93–94), Pearl (2000, pp. 349–350). Note that I am not claiming that it doesn't make sense to talk about an intervention on the whole universe—and therefore that causation only applies to small worlds on the interventionist picture—just that interventionist theories apply more naturally to small worlds.

²⁵ Of course, there is a sense in which the causal structure remains the same throughout the scenario described above: the laws of nature don't change, it is just that Hooke's law has limited application. It would be possible to use a lower level causal model to describe both the case when the spring breaks and when it doesn't. The point is that we do use *ceteris paribus* laws (including Hooke's law) as a guide to causal reasoning.

or thermodynamic control, and they use this ability to control whether the kinetic or thermodynamic product is isolated.

Our ability to manipulate which causal system is instantiated is not restricted to organic chemistry. Consider a transformer: a toy that can be converted from a car into a robot (and back again). The process of transforming the toy between these two states can be thought of as an intervention on the causal system instantiated, and thus the causal structure of the toy. For example, placing the transformer at the top of an inclined plane will cause it to roll down the plane when it is a car, but not when it is a robot. Playing with a transformer is therefore another situation in which we intervene on the causal system. This is actually a very common phenomenon, and one that greatly extends our ability to manipulate the world.

The connection between causation and manipulation that is at the heart of Woodward's interventionism makes manipulations to the causal system especially important for this theory. Recall that according to Woodward, the point of having a concept of causation is to enable us to identify correlations that are exploitable for the purpose of manipulation or control. The main goal of organic chemistry—a goal that has been extremely successful—is to synthesise new compounds (or existing compounds in new ways). In Goodwin's words, the goal of this discipline is therefore really 'to assess and expand the human capacity to manipulate and control' the domain of organic molecules (2012, pp. 445–446). I have shown that this capacity to manipulate and control includes the ability to manipulate which system is instantiated, as well as to intervene on the values of variables within a causal system. This former ability allows chemists, not just to *identify* correlations that are exploitable for the purpose of manipulation and control, but to *create* them.

Since the development of the capacity to control our environment is precisely the point of having a concept of causation in the first place, according to Woodward, the tools that chemists use to achieve this mastery are highly relevant to his theory. Interventions on the causal system therefore represent an important addition to his theory, and further work needs to be done to formalise this kind of intervention.

8 Interventionism and pragmatism

The discussion in this paper has implications for another issue that is central to interventionism. It is often pointed out that there is more than one way of modelling any given situation, some of which we think are better than others.²⁶ This raises the question: what does it take to be an appropriate causal model? There is general agreement that which variables are included in a causal model is partly a matter of pragmatics.²⁷ In Woodward's terminology, the variables included in a causal model are those that represent 'serious possibilities' (2003, pp. 86–91). For example, when modelling the cause of a forest fire, we may not include the variable {amount of oxygen in the atmosphere}, even though the occurrence of a fire is counterfactually dependent on the presence of oxygen. According to Woodward, this is because changes in the oxygen

²⁶ For example, see [Strevens \(2007, 2008\)](#); [Woodward \(2008\)](#).

²⁷ See [Joseph and Christopher \(2010\)](#), [Blanchard and Schaffer \(2015\)](#) and [Woodward \(2003\)](#).

content of the atmosphere do not represent a serious possibility. Since which changes to actuality are serious possibilities is highly context sensitive, the notion of a serious possibility admits some degree of interest-relativity into the concept of causation.

In a recent paper, Woodward addresses a very closely related question: the ‘problem of variable choice’ (2016). He starts with the assumption that some choices of variables are better than others for the purposes of causal analysis, and then considers which criteria we should take into account when deciding which variables to include in a causal model. When framing this discussion, he acknowledges that the appropriate choice of variables is determined partly by the goal of an enquiry:

[T]he problem of variable choice should be approached within a means/ends framework: cognitive inquiries can have various goals or ends and one can justify or rationalize candidate criteria for variable choice by showing that they are effective means to these ends. (2016, p. 1051)

However, Woodward maintains that there is only one goal that is relevant to the construction of causal models, namely causal representation/explanation:

[B]y goals of inquiry, I have in mind broad cognitive aims like provision of causal explanations as opposed to, say, “prediction”, rather than the idiosyncratic preferences of particular investigators for certain variables (such as the preferences of some psychologists for only behavioural variables). (2016, Fn. 7)

Thus, Woodward takes the construction of an appropriate causal model to be a matter of pragmatics in two quite limited senses. First, the context determines which variables represent serious possibilities; second, some choices of variables are better than others given the overall goal of causal representation.

In this paper, I have discussed several causal models. These are influenced by the theory of organic chemistry in two ways. First, the variables included in each model are those that organic chemists consider relevant—changes to the values of these variables are ‘serious possibilities’. Second, the partitioning of reactions into those under kinetic, thermodynamic, and kinetic versus thermodynamic manipulation is based on organic theory. I now show that this second sense in which the models are influenced by the theory of organic chemistry leads to a more extensive pragmatism than Woodward has acknowledged.

8.1 The pragmatic nature of organic chemistry

Goodwin has shown that the theory of organic chemistry has been heavily influenced by its major goal, the synthesis of novel compounds (2012, 2013). He focuses on explanations of the outcomes of reactions that cite structural features of the starting materials (see Sect. 3.1 of this paper). However, the categorisation of reactions into those under kinetic, thermodynamic, and kinetic versus thermodynamic manipulation also aids the goal of synthesis. For one thing, when a reaction is understood to be under one of these forms of manipulation, it is possible to work out how to alter the reaction conditions in order to maximise the yield of the desired product.

Comparing the theory of organic chemistry with another way of representing organic reactions makes it easier to see how the former has been influenced by the goal of synthesis. It is possible to computationally determine a potential energy surface for any given reaction using quantum mechanical calculations. A potential energy surface is a three dimensional version of a reaction profile diagram (see Fig. 1), which plots energy against different atomic geometries. These diagrams (or the quantum mechanical calculations that they are based on) can be used to predict the outcome of a reaction far more accurately than the kinetic and thermodynamic methods discussed previously. For example, it is possible to determine the energy of the transition state from a potential energy surface, and therefore calculate the rate of reaction. It is also possible to determine the difference in energy between reactants and products, and therefore the free energy of the reaction. It may therefore be tempting to say that the kinetic and thermodynamic models discussed above can be reduced to potential energy surfaces, and that chemists would be better off thinking about reactions exclusively in these terms. There are two reasons for thinking that organic chemistry cannot be replaced by quantum chemistry, however.

First, thinking about reactions as being either kinetic or thermodynamic control, and therefore as governed by either kinetic or thermodynamic laws, respectively, is far simpler than calculating a potential energy surface, because the kinetic and thermodynamic laws are far more computationally tractable than the quantum mechanical calculations. There are therefore practical reasons for using simple kinetic and thermodynamic models for making predictions and giving explanations in organic chemistry.

Second, and more importantly, Goodwin shows that quantum chemistry and organic chemistry have different ‘explanatory agendas’, which means that they appeal to different generalisations when making predictions and giving explanations. Effectively, they provide information about different possibility spaces (2013, p. 1161). This means that it is possible to make predictions on the basis of organic theory that can’t be derived from quantum chemistry alone.

Though the quantum chemist might be able to provide a detailed potential energy surface for the reaction in question ... this does not by itself tell the chemist what structural features might be changed in order to modify the outcome. (Goodwin 2013, p. 1166)

Goodwin’s point is that although quantum chemistry allows you to create a potential energy surface (from which you can determine the activation energy and the free energy of possible reaction pathways), this potential energy surface doesn’t tell you anything at all about related reactions. If you are an organic chemist trying to synthesise a product that has never been made before, you need to be able to predict the outcome of novel reactions, based on similar, known reactions. This is exactly what the theory of organic chemistry—including reaction kinetics and thermodynamics—is designed for, but it is something that quantum chemistry, by itself, is unable to do. Thus, the theory of organic chemistry can’t simply be replaced by quantum chemistry.

As Goodwin points out, the claim that different scientific disciplines have different explanatory agendas (and that these result in different choices of variables) is one of

the features of Woodward's interventionism (2013, pp. 1160–1161).²⁸ However, this raises the question: what counts as an acceptable explanatory agenda?

We have seen that in his most recent work, Woodward claims that the only goal that constrains which variables should be included in the causal model of a particular system is that of causal representation (or causal explanation). The theory of organic chemistry, however, is not simply guided by the goal of causal representation. It is heavily influenced by a more practical goal: the desire to develop new syntheses. The concepts *kinetic control* and *thermodynamic control* were developed in this context: outside a synthetic context, these concepts—and the corresponding variables {identity of the kinetic product} and {identity of the thermodynamic product}—wouldn't have originated at all.

The range of acceptable explanatory agendas therefore includes more than just 'causal representation', but also the capacity to manipulate a domain. I have also shown that which interpretation of kinetic versus thermodynamic manipulation is appropriate depends on a scientist's explanatory agenda. More generally, what makes a variable set (and therefore a causal model) appropriate is determined partly by the particular goal of a causal or explanatory enquiry, not solely the broad cognitive aim of causal representation and the causal structure of the relevant part of the world. The pragmatic elements of interventionist causal modelling therefore extend beyond those that Woodward himself endorses—the fact that different causal enquiries generate different causal models suggests that interventionism leads to a version of pragmatism about causation.

The claim that interventionism is associated with a form of pragmatism about causation has also been made with respect to psychiatry. K. S. Kendler and J. Campbell discuss the application of interventionist causal models to this field. They argue that interventionism 'can provide a single, clear empirical framework for the evaluation of all causal claims of relevance to psychiatry' (2009, p. 881), while avoiding metaphysical issues that tend to 'preoccupy' psychiatrists. What they are claiming is that thinking about causation in interventionist terms allows psychiatrists to focus on the important practical issue—how best to prevent and treat mental illness—without getting caught up in metaphysical distractions. Longino points out that this emphasis on practical solutions at the expense of metaphysics involves embracing a form of pragmatism (2013, pp. 149–150).

Like organic chemistry, psychiatry is a practical field, in which the main aim is to manipulate parts of the world in order to achieve certain goals. These investigations into the application of interventionism to organic chemistry and psychiatry suggest that interventionism supports pragmatism about causation: a pragmatism that helps us to achieve practical goals.

9 Conclusion

I have shown that kinetic, thermodynamic, and kinetic versus thermodynamic manipulation can all be represented using interventionist causal models, and therefore that

²⁸ See Woodward (2003, pp. 231–233).

the associated causal claims made by organic chemists can be accounted for by Woodward's interventionism. However, I showed that kinetic versus thermodynamic manipulation can also be thought of as involving interventions on the causal system itself. In an interventionist context it makes sense—and is indeed necessary—to talk about interventions that make a difference to which causal system is instantiated. This situation therefore exemplifies a kind of causal reasoning that has not previously been discussed, to which interventionists, in particular, need to pay more attention. Finally, I argued that the application of interventionism to organic chemistry has a more general implication: it suggests that interventionism leads to pragmatism about causation.

Acknowledgments I would like to thank Huw Price, Hasok Chang, Claire Benn, Alison Fernandes, Shyane Siriwardena, Mat Simpson, two anonymous referees for *Synthese*, and participants at the Annual International Society for the Philosophy of Chemistry Conference for providing valuable feedback on various drafts of this paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Ackermann, L., El Tom, D., & Fürstner, A. (2000). Ruthenium carbene complexes with imidazol-2-ylidene ligands: Syntheses of conduritol derivatives reveals superior RCM activity. *Tetrahedron*, 56, 2195–2202.
- Blanchard, T., & Schaffer, J. (2015). Cause without default. In H. Beebe, C. Hitchcock & H. Price (Eds.), *Making a difference: Essays in honour of Peter Menzies*. Oxford: Oxford University Press.
- Cartwright, N. (1979). Causal laws and effective strategies. *Noûs*, 13, 419–437.
- Clayden, J., Greeves, N., & Warren, S. (2012). *Organic chemistry* (2nd ed.). Oxford: Oxford University Press.
- Goodwin, W. (2003). Explanation in organic chemistry. *Annals of the New York Academy of Sciences*, 988, 141–153.
- Goodwin, W. (2012). Experiments and theory in the preparative sciences. *Philosophy of Science*, 79, 429–447.
- Goodwin, W. (2013). Quantum chemistry and organic theory. *Philosophy of Science*, 80, 1159–1169.
- Hitchcock, C. (1996). Farewell to binary causation. *Canadian Journal of Philosophy*, 26, 267–282.
- Hitchcock, C. (2001). The intransitivity of causation revealed in equations and graphs. *The Journal of Philosophy*, 98, 273–299.
- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *Philosophical Review*, 116, 495–532.
- Joseph, H., & Christopher, H. (2010). Actual causation and the art of modelling. In R. Dechter, H. Geffner, & J. Y. Halpern (Eds.), *Probability and causality: A tribute to Judea Pearl* (pp. 383–406). London: College Publications.
- Kendler, K. S., & Campbell, J. (2009). Interventionist causal models in psychiatry: Repositioning the mind-body problem. *Psychological Medicine*, 39, 881–887.
- Lewis, D. (1986). *Philosophical papers* (Vol. 2, pp. 241–269). Oxford: Oxford University Press.
- Longino, H. E. (2013). *Studying human behaviour: How Scientists investigate aggression and sexuality*. Chicago, IL: University of Chicago Press.
- Maslen, C. (2004). Causes, contrasts, and the nontransitivity of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 341–358). Cambridge, MA: MIT Press.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.
- Ramsey, J. L. (2008). Mechanisms and their explanatory challenges in organic chemistry. *Philosophy of Science*, 75, 970–982.
- Schaffer, J. (2005). Contrastive causation. *Philosophical Review*, 114, 297–328.

- Schaffer, J. (2013). Causal contextualism. In M. Blaauw (Ed.), *Contrastivism in philosophy* (pp. 33–63). New York: Routledge.
- Statham, Georgie. (2016). Contrastive causal claims: A case study. *British Journal for the Philosophy of Science* (Published online).
- Strevens, M. (2007). Review of Woodward, 'Making Things Happen'. *Philosophy and Phenomenological Research*, 74, 233–249.
- Strevens, M. (2008). Comments on Woodward, 'Making Things Happen'. *Philosophy and Phenomenological Research*, 77, 171–192.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.
- Woodward, J. (2007). Causation with a human face. In H. Price & R. Corry (Eds.), *Causation, physics, and the constitution of reality* (pp. 66–105). Oxford: Oxford University Press.
- Woodward, J. (2008). Response to Strevens. *Philosophy and Phenomenological Research*, 77, 193–212.
- Woodward, J. (2016). The problem of variable choice. *Synthese*, 193, 1047–1072.